

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
31 January 2002 (31.01.2002)

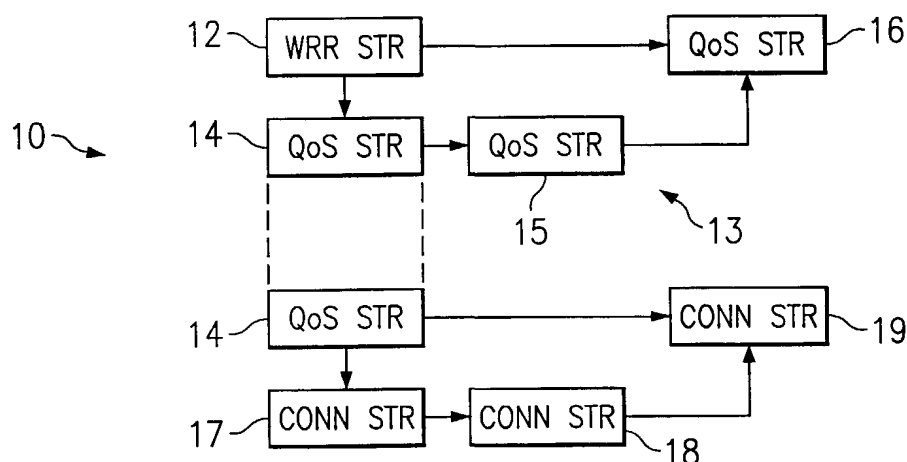
PCT

(10) International Publication Number  
**WO 02/09358 A2**

- (51) International Patent Classification<sup>7</sup>: **H04L 12/00**
- (21) International Application Number: PCT/US01/23474
- (22) International Filing Date: 25 July 2001 (25.07.2001)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
09/626,400 26 July 2000 (26.07.2000) US
- (71) Applicant (for all designated States except US): **SAN-TERA SYSTEMS, INC.** [US/US]; 2901 Summit Avenue, Suite 100, Plano, TX 75074 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **LI, Na** [CN/US]; 11316 Jollyville Rd., #171, Austin, TX 78759 (US). **LI, San-qi** [US/US]; 5508 Bassinghall Lane, Plano, TX 75093 (US).
- (74) Agents: **JEANG, Wei Wei** et al.; Munsch Hardt Kopf & Harr, P.C., 4000 Fountain Place, 1445 Ross Avenue, Dallas, TX 75202-2790 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**  
— without international search report and to be republished upon receipt of that report

[Continued on next page]

(54) Title: METHOD OF ACTIVE DYNAMIC RESOURCE ASSIGNMENT IN A TELECOMMUNICATIONS NETWORK



(57) Abstract: A method of adaptive resource allocation in transmitting data is provided. In particular, the method of the present invention is applicable to scheduling bandwidths using a modified and dynamic weighted round robin process. The method includes the steps of allocating a resource to each of a plurality of data transmitting active connections, the plurality of active connections belonging to more than one class of service. Then the method determines a lender class of service for each active connection from which resources may be reallocated to the active connection, and periodically comparing the resource usage of an active connection to upper and lower thresholds. Borrowing action for resources and returning action of borrowed resources are carried out in response to the current usage exceeding the upper threshold or less than the lower threshold, respectively.



WO 02/09358 A2



---

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

## METHOD OF ACTIVE DYNAMIC RESOURCE ASSIGNMENT IN A TELECOMMUNICATIONS NETWORK

### TECHNICAL FIELD OF THE INVENTION

This invention relates in general to telecommunications network and equipment, and more particularly to a method of active dynamic resource assignment in a telecommunications network.

### 5 BACKGROUND OF THE INVENTION

When a new connection is to be set up by a switch in a telecommunications network, the switch has to assign certain resources, i.e., buffer to hold the incoming cells and bandwidth to transmit cells, according to the specified Quality of Service (QoS) requirement and the connection parameters. QoS can be specified in terms of the amount of delay and/or cell loss ratio. The connection parameters include  
10 average sending rate, peak sending rate, and maximum back-to-back cell number (maximum burst size). Initially, the resource assignment is static and remains unchanged during the lifetime of the connection. Therefore, if more resources become available, a static resource connection assignment cannot make use of the additional available resources to achieve better performance.

In some conventional switches, dynamic adjustment of the buffer (memory) allocation for  
15 connections based on the overall queue occupancy is performed. When most connections are active, and the total buffer occupancy is high, the per connection buffer allocation is low. When the total buffer occupancy is low, the per connection buffer assignment is high, permitting active connections using more buffer to achieve lower cell loss rate.

The same requirement is true for bandwidth assignment. There are three conventional scheduling  
20 methods, namely strict priority, rate scheduling, and weighted round robin. Strict priority scheduling is the simplest method. Strict priority scheduling always gives connections in one class higher priority than connections in another class. For example, whenever any connection in class A has cells to transmit, cells of connections in class B are not sent. Strict priority is risky since there is no control over how much of the class B connections will be made. Therefore, it is difficult to support more than two classes of  
25 service. Furthermore, service guarantee can only be made for the highest priority class, and not to the others. Strict priority also cannot be used for scheduling allocation within one class.

The rate scheduling method uses a timed ring to make the cell-to-cell spacing of a connection equal to a predefined value, which is the inverse of the connection's assigned bandwidth. When the assigned bandwidth is static, the cells associated with the connection cannot be sent faster than the  
30 assigned bandwidth.

Weighted round robin (WRR) is another conventional scheduling method. It uses the weight assigned to a connection to make the scheduling decision. The actual bandwidth allocated to a connection

is proportional to its own weight and inversely proportional to the total weight of all active connections. Therefore, when other connections are not using their bandwidths, i.e. they have no cells to transmit, a connection can automatically get more bandwidth. WRR scheduling can provide a connection with the assigned bandwidth and make a full use of the available bandwidth as long as the switch bounds the total weight of active connections by rejecting some connections. For example, if there are a total of ten connections and each with the weight of 2. If all connections are active, a connection is allocated 1/10 link bandwidth. If only five connections are active, then an active connection is now allocated 1/5 link bandwidth. However, this method uses static weight assignment, which only occurs at connection setup. Given a bounded total weight, the weight assigned to a connection with a delay guarantee has to be large enough to make sure that the connection obtains sufficient bandwidth without violating the QoS service contract. If the weight assignment is conservative and based on the largest possible sending rate, the switch can obey the service contract but only support a very limited number of connections that require strict QoS guarantee. On the other hand, if the weight is assigned based on the average sending rate, it is possible that the cells experience large delay and are useless to the receiver. Internet service providers (ISPs) prefer to assign minimum weight to a connection in order to support more connections for higher profit and still adhere to the service contract. WRR with static weight assignment cannot fulfill these goals.

#### **SUMMARY OF THE INVENTION**

In accordance with the present invention, a dynamic adaptive resource allocation method is provided which overcomes problems associated prior methods. During the life of a connection, the weight assignment of the connection is adaptively adjusted according to its dynamic bandwidth or rate requirement. The bandwidth or rate is borrowed from a lender class of services, not from a specific connection, and is returned within a maximum borrow time or when it is no longer needed. Therefore, delay-critical connections starved for bandwidth can borrow resources to adhere to its QoS requirement and return it when the borrowed resource is no longer needed.

In one embodiment of the present invention, a method of adaptive resource allocation in transmitting data is provided. In particular, the method of the present invention is applicable to scheduling bandwidths using a modified and dynamic weighted round robin process. The method includes the steps of allocating a resource to each of a plurality of data transmitting active connections, the plurality of active connections belonging to more than one class of service. Then the method determines a lender class of service for each active connection from which resources may be reallocated to the active connection, and periodically comparing the resource usage of an active connection to an upper threshold and a lower threshold. Borrowing action for resources from the lender class of service for the active connection is carried out in response to the current usage exceeding the upper threshold, and

returning action for borrowed resources to the lender class of service from the active connection in response to the current usage being less than the lower threshold is carried out.

In another embodiment of the present invention, a telecommunications equipment having limited resources to allocate to active connections for a plurality of classes of services, using a method of dynamic resource allocation, includes a first data structure storing resource allocation weighting data associated with each class of services, and a second data structure storing resource allocation weighting data associated with each active connection in each class of services. The second data structure includes a class identifier specifying a lender class of services from which resource may be borrowed, an upper threshold indicative of when resource borrowing action for this active connection should be initiated,

a lower threshold indicative of when resource returning action for this active connection should be initiated, and a sequence of data associated with each resource borrowing action of resources that have not been returned.

In yet another embodiment of the present invention, a method of adaptive dynamic weight assignment for transmitting data includes the steps of assigning a total weight to each class of services, assigning a weight to each of a plurality of data transmitting active connections in each class of services, and determining a lender class of service for each active connection from which weighting may be reallocated to the active connection. The method periodically compares the resource usage of an active connection to an upper threshold and a lower threshold, and borrows additional weighting from the lender class of service for the active connection in response to the current usage exceeding the upper threshold. The borrowed weighting is returned to the lender class of service from the active connection in response to the current usage being less than the lower threshold.

#### **BRIEF DESCRIPTION OF THE DRAWINGS**

For a more complete understanding of the present invention, the objects and advantages thereof, reference is now made to the following descriptions taken in connection with the accompanying drawings in which:

FIGURE 1 is a block diagram of a weighted round robin scheduling data structure;

FIGURE 2 is a block diagram of WRR information data structure;

FIGURES 3A, 3B and 3C are a block diagrams of an embodiment of data structures used to store adjusted weight data according to the teachings of the present invention;

FIGURE 4 is a flowchart of an embodiment of a method of adjusting connection weighting according to the teachings of the present invention;

FIGURE 5 is a flowchart of an embodiment of a method of borrowing bandwidth according to the teachings of the present invention;

FIGURE 6 is a flowchart of an embodiment of a method of returning bandwidth according to the teachings of the present invention;

FIGURE 7 is a flowchart of an embodiment of a method of updating the data record after borrowing bandwidth according to the teachings of the present invention; and

5 FIGURE 8 is a flowchart of an embodiment of a method of updating the data record after returning bandwidth according to the teachings of the present invention.

#### DETAILED DESCRIPTION OF THE DRAWINGS

10 The preferred embodiment of the present invention and its advantages are best understood by referring to FIGURES 1 through 8 of the drawings, where like numerals are used to refer to like and corresponding parts of the various drawings. To facilitate the understanding of the following description, acronyms and variable names are listed below for ease of reference:

$\Delta BW$	amount of bandwidth the connection can borrow each time
$\Delta t$	minimum time requirement between borrow and/or return actions
$\Delta W\_F$	the weight increment for the connection
$\Delta W\_QoS$	the reduced weight amount of the lender class of services
AWD	adjust weight descriptor
AWT	adjust weight table
Borrow_time_limit	maximum time to return borrow bandwidth
BW	bandwidth assigned to a scheduler
CONN ID	connection identifier
FTP	file transfer protocol
ISP	Internet service provider
q	buffer occupancy of a connection
q_down	buffer lower threshold
q_up	buffer upper threshold
QoS	quality of service
QoS_ID	class identifier that a connection can borrow bandwidth from
t_last	the last time a borrow or return action took place for the connection
timestamp	time of borrow action stored in adjust weight descriptor
W	total weight of all QoS classes
W_conn	weight of a connection stored in the Conn data structure
W_QoS	weight of a class associated with a connection stored in the QoS data structure
W_QoS_b	weight of the class associated with the current connection borrowing bandwidth

W_QoS_F	total weight of connections belonging to a class of connections
W_QoS_l	weight of the class of service to lend bandwidth
WRR	weighted round robin

FIGURE 1 is a simplified block diagram of a weighted round robin (WRR) scheduling data structure 10. WRR structure (WRR str) 12 is a data structure that stores the parameters associated to a particular WRR scheduler. WRR str 12 maintains a linked list 13 of the QoS classes data structure (QoS str) 14-16 and possibly the total weight of all QoS classes that belong to this scheduler. Each QoS class has its own weight, and maintains its own linked list of connection data structure, Conn str 17-19. QoS str 14 keeps such information, e.g., the head and tail pointer of the link list of connections in this class, also possibly the total weight of such connections. Each connection has its own weight as well, which is stored in Conn str 17-19. Only active connections are linked in the scheduler, therefore the bandwidth is only shared among active connections.

In WRR, the bandwidth of a flow  $j$  can be expressed as:

$$BW_{conn_j} = (W_{conn_j} / \text{sum}(W_{conn_i})) * (W_{QoS_m} / \text{sum}(W_{QoS_n})) * BW, \quad (1)$$

Where  $BW$  is the bandwidth assigned to the WRR scheduler;  $W_{conn_j}$  is the  $j$ th connection's weight; and  $W_{QoS_m}$  is the weight of QoS class  $m$ , assuming connection  $j$  is in class  $m$ . When the weight of a connection is increased, its bandwidth is increased accordingly.

FIGURE 2 is a block diagram of various data fields involved in the methodology. The incoming cell header data structure 10 includes a connection identifier (CONN ID) 22. Conn str includes a buffer occupancy  $q$  24 of a connection, and a weighting for the connection  $W_{conn}$  26.  $W\_QoS\_F$  field 28 is the total weight of connections belonging to a class, and  $W\_QoS$  30 is the weight for a class, both stored in the QoS structure.  $W$  32 is the total weight of all classes and  $BW$  34 is the allocated bandwidth stored in a WRR structure 12. CONN ID 22,  $q$  24,  $W_{conn}$  26, and  $W\_QoS$  30 are conventional data fields for standard switches that support per connection QoS.  $W\_QoS\_F$  28 and  $W$  32 are available in some switches, but not in others. In conventional switches,  $BW$  34 is a fixed value at scheduler initialization.

FIGURES 3A, 3B and 3C are block diagrams of an embodiment of data structures used to store adjusted weight data according to the teachings of the present invention. An Adjust\_Weight\_Table (AWT) 40 includes multiple AWT entries 41, with each AWT entry 41 pointing to or containing the address of the head 43 and tail 45 of a linked list 42 of Adjust\_Weight\_Descriptors (AWD) 43-45. Each AWT entry 41 corresponds to an established connection. CONN ID 22 in cell header 20 is used to index

or address the AWT entry 41. AWD 43 is a data structure that stores the predefined parameters and dynamic variables for the connections. AWD records the information of each bandwidth borrowing action. A connection can borrow bandwidth for several times before returning any bandwidth.

In an AWT entry 41, a head\_pointer 50 points to the address of the AWD that is at the head 43 of link list 42, and tail\_pointer 52 points to the address of the AWD 45 that is at the tail of link list 42. t\_last 54 is a data field that records the timestamp of the latest update action. The update action can be borrowing or returning bandwidth. These three variables 50-54 are dynamically updated. Weight\_adjust\_enable 56 indicates if this connection has authorization to borrow bandwidth from other connections. Weight\_adjust\_enable 56 is initialized during connection setup. If this flag is set, the following fields should be initialized as well. Qos\_ID 58 is the class identifier that this connection can borrow bandwidth from. q\_up 60 is the buffer threshold used to determine if the connection should borrow bandwidth. q\_down 62 is the buffer threshold used to determine if the connection should return borrowed bandwidth.  $\Delta t$  64 is the time threshold used to limit how often the update action can occur.  $\Delta BW$  66 is the amount of bandwidth this connection can borrow each time. Borrow\_time\_limit 68 may be an optional field which defines how long the connection can borrow bandwidth and keep it before returning it. This check is important since it is possible that the source of the borrowing connection sends more data than permitted in the contract or that the bandwidth initially allocation to the connection is not reasonable. If the source sends too much data, the User Parameter Control (UPC) at the switch entrance of the connection can be configured to discard excess cells to prevent a connection from maliciously benefiting from this invention when resource is stringent, and let the cells pass when resource is free. If the resource allocation is not reasonable, this problem can be identified before the connection suffers and experiences performance degradation. Action code 70 defines what action to take during abnormal state. In this example, the abnormal state can be defined as when a connection borrows bandwidth longer than the Borrow\_time\_limit. The abnormal state can be defined in other ways.

In AWD 43-45, timestamp 72 records the time when a borrow action occurs.  $\Delta W_{QoS}$  74 is the reduced weight amount of the borrowed class.  $\Delta W_F$  76 is the increased weight amount of this connection. A next field 78 points to the next AWD in link list 42. Whenever an eligible connection needs to borrow bandwidth from a class, memory is dynamically created and allocated to hold a new AWD. Proper values are filled into this new AWD, which is then linked into the tail 45 of linked list 42. The detailed procedure is described below. When the connection needs to return the borrowed bandwidth, the information of the AWD at the head of link list 42 is read for weight updates. The head AWD 43 is then removed from linked list 42, and the associated memory is released using standard procedures.



The switch (not shown) also uses a flag, ADBA (adaptive dynamic bandwidth assignment) enable 80, to indicate if the present invention is enabled or disabled. This parameter is set when the switch is initialized, and can be modified dynamically by the administrator.

FIGURE 4 is a flowchart of an embodiment of a method of dynamic connection weight adjustment 90 according to the teachings of the present invention. In this example, a determination of a predetermined event - a cell arrival - is made, as shown in block 92. Alternatively, periodic determination of whether connection weights should be adjusted can be initiated. At a cell arrival, a decision is made whether to adjust the weights. If adjustment is not permitted, as determined by the value of ADBA enable flag 80 in block 94, normal procedure or treatment is followed. If ADBA enable 80 is set, the connection ID (CONN ID) 22 of the incoming cell is read from memory, as shown in block 98. Using CONN ID 22, the value of weight\_adjust\_enable 56 is accessed from AWT entry 41, as shown in block 100. If weight\_adjust\_enable flag 56 is not set, execution stops in block 102. Otherwise,  $t_{last}$  54 (the latest borrow action time),  $\Delta t$  64 (the shortest time allowable between borrow actions), and borrow\_time\_limit 68 (how soon borrowed bandwidth needs to be returned) are read from AWT entry 41, as shown in block 104. If the current time,  $t_{now}$ , minus  $t_{last}$  54 is smaller than  $\Delta t$ , as determined in block 106, then the process ends. Therefore, update action can only occur at most once in each  $\Delta t$  time interval. If this threshold is satisfied, timestamp 72 of AWD 43 at the head of linked list 42 of the connection maintained in AWT entry 41 is obtained, as shown in block 108. If the current time minus timestamp 72 is larger than borrow\_time\_limit 68, as determined in block 110, the connection is requested to return some borrowed bandwidth by assuming the action code is returning the bandwidth 112. Return bandwidth process 112 is described in more detail in FIGURE 6. Other actions can also be defined and initiated. Otherwise, the current buffer occupancy of the connection  $q$  24 is read from the switch memory and compared with  $q_{down}$  62 from AWT entry 41 in block 114. If  $q < q_{down}$ , the connection returns the bandwidth. Otherwise, if  $q > q_{up}$ , as determined in block 116, the connection borrows bandwidth 118, which is described in more detail in FIGURE 5. If  $q$  is not greater than  $q_{up}$ , execution stops.

The value of  $q_{up}$  60 depends on the delay requirement of the connection and the average bandwidth allocated to the connection. The customer delay requirement is end-to-end. The delay consists of propagation delay, processing delay, and queuing delay. The propagation delay and processing delay are first subtracted from the delay parameter since they are fixed value. Next, the end-to-end value is broken into local values, i.e., each switch along the connection path is assigned its local queuing delay. The multiplication of the (local queuing delay -  $\Delta t$ ) and the average bandwidth allocated to the connection gives a reasonable value for  $q_{up}$ . More delicate tuning is possible. When  $q_{down}$  is smaller than  $q_{up}$ , it allows hysteresis for oscillation avoidance. A reasonable value for  $q_{down}$  is half of the  $q_{up}$ .

FIGURE 5 is a flowchart of an embodiment of a method of borrowing bandwidth 118 according to the teachings of the present invention. The related parameters are read and  $\Delta W\_QoS$  74 and  $\Delta W\_F$  76 are calculated, as shown in blocks 120 and 122.  $W\_QoS$  is the weight of the class associated with the current connection.  $W$  32 is the total weight of the classes.  $BW$  is the bandwidth of the WRR scheduler.  $W\_QoS\_F$  is the total weight of connections in this class. It then compares the weight of the class to lend the bandwidth  $W\_QoS\_l$  with  $\Delta W\_QoS$ , as shown in block 126. The class to lend the bandwidth is identified by  $QoS\_ID$  in the AWT entry. If  $W\_QoS\_l < \Delta W\_QoS$ , ADBA stops in block 128 since there is insufficient bandwidth to borrow from. This situation can optionally trigger an alarm. Otherwise, the  $W\_QoS\_l$  is subtracted by  $\Delta W\_QoS$  in block 130. The weight of the class that associated with the current connection  $W\_QoS\_b$  is added by  $\Delta W\_QoS$  in block 132. Also, the connection weight  $W\_conn$  is added by  $\Delta W\_F$  in block 124. The AWD records are then updated in block 136, which is described in more detail in FIGURE 7.

FIGURE 7 is a flowchart of an embodiment of a method 136 of updating the data record after borrowing bandwidth according to the teachings of the present invention. A new AWD is allocated and filled with the current time (timestamp),  $\Delta W\_QoS$  and  $\Delta W\_F$ , as shown in block 138. The AWD is then linked to the tail 45 of the link list 42 so that it now is the new tail of the linked list, as shown in block 140.  $t\_last$  54 in the corresponding AWT entry 41 is also updated with the current time, as shown in block 142. Optionally, an action log file kept for documentation purposes is updated to add the latest borrow action information, as shown in block 144. The process ends in block 146.

According to the teachings of the present invention, the bandwidth is borrowed from a predefined class, not from specific connections. This approach reduces the impact on individual connections. On the other hand, all borrowed bandwidth is assigned to a particular connection since only this particular connection needs the extra bandwidth.

FIGURE 6 is a flowchart of an embodiment of a method 112 of returning bandwidth according to the teachings of the present invention. In blocks 150 and 152,  $\Delta W\_QoS$  74 and  $\Delta W\_F$  76 are read from AWD 43 at the head of the link list 42 maintained by the AWT entry 41. The  $W\_QoS\_l$  identified by  $QoS\_ID$  58 in AWT entry 41 is read and increased by  $\Delta W\_QoS$  (blocks 154 and 156).  $W\_QoS\_b$  is then decreased by  $\Delta W\_QoS$  in block 158, and the  $W\_conn$  is decreased by  $\Delta W\_F$  in block 160. The borrowed bandwidth is thus returned. The AWD record is updated in block 162, which is described in more detail in FIGURE 8.

FIGURE 8 is a flowchart of an embodiment of a method 162 of updating the data record after returning bandwidth according to the teachings of the present invention. AWD 43 at the head of linked list 42 is removed from the linked list and deleted, as shown in block 164. The memory allocated to the deleted AWD may be garbage collected and returned to the available memory pool. In block 166,  $t\_last$

54 is updated with the current time. In block 168, the action is then logged into the action log file, if this documentation is maintained. The process ends in block 170.

The present invention permits initially assigning weight to a connection that requires delay guarantee only based on its average sending rate. During the lifetime of the connection, it then actively  
5 adjusts the weight of the connection to match its dynamic bandwidth requirement. Therefore, an ISP can support more primary connections while obeying the service contracts. For the same performance, the customers are likely to be charged less than with the current conventional technology. This provides a strong competitive advantage to an ISP who adopts this invention.

There are different types of connections in a switch. Some requires very short instantaneous  
10 delay, like voice connections and videoconferencing connections. Some others are less sensitive to the instantaneous delay, like telnet. Connections like file transfer protocol (FTP) do not care about the instantaneous delay as long as the average delay over the lifetime of the connection is reasonable. By adjusting the weights, the invention actively borrows bandwidth from delay non-sensitive connections for a delay-critical connection when the connection is going to break its limit. When the delay-critical  
15 connection no longer needs the borrowed bandwidth or because of other policies, the borrowed bandwidth is returned.

Allowing the return of bandwidth is important for many reasons. The maximum weight is constrained by hardware, and cannot be increased infinitely. Using the present invention, only urgent delay-critical connections that will exceed the respective delay requirements are allowed to borrow  
20 bandwidth. Other connections that meet the respective delay budget can still use their original weight without alterations. Therefore, the same bandwidth can be borrowed by one connection, returned, and then borrowed by another delay-critical connection. By getting borrowed bandwidth back, the lender connections can also have small average delay.

There are two bandwidth return procedures. One is explicit, as described in FIGURE 6. Another  
25 one is implicit, when the arriving rate of a borrowing connection is below the average rate, the unused bandwidth is automatically returned to other connections, including the lender connections, through the original WRR scheduling. With the original WRR, bandwidth allocation based on the average rate is not sufficient for delay-sensitive connection with variable bandwidth requirement. By introducing the explicit borrowing and explicit returning bandwidth of the invention, the delay requirement can be satisfied while  
30 allocating bandwidth according to the average rate. The implicit bandwidth is still necessary even with the explicit bandwidth return. The reason is that lending some bandwidth and getting it back later is not the same as never lending the bandwidth. Actually, the former action results in less data being transmitted by the lender connections. The compensation is done through the implicit bandwidth return. The invention naturally supports multiple classes with different bandwidth borrowing and return policies.  
35 The variant QoS requirement can be satisfied at different classes. The conventional WRR with fixed

weight does not prioritize the QoS requirement and importance of the different connections, while strict priority goes to the extreme of the priority. Strict priority is risky since there is no control on how much the low priority connections can be starved. The service guarantee can only be committed to the highest priority class, not to the others. On the contrary, the present invention prioritizes the connections and performs the bandwidth assignment in a tightly controlled and very flexible manner.

The connections can be divided into several classes according to their QoS requirement. Classifying a connection into which class is very flexible, depending on the policy configured by the ISP. The eligibility of borrowing and lending is also very flexible, policy dependent, can be based on the class or based on per-connection, or even based on which class currently has most bandwidth to lend etc. The maximum amount of bandwidth and how long the bandwidth a connection can borrow without returning are also configurable. For example, a voice connection can be classified as class A, a FTP connection can be classified as class B. Connections in class A can borrow bandwidth from class B, but not vice versa. The configuration may also specify that connections in class A can only borrow bandwidth for one hour without return. Otherwise, it has to return some bandwidth. In addition, the borrowing and returning actions can be logged for billing, traffic analysis, and system planning purposes.

Although the present invention has been described in terms of active dynamic bandwidth assignment by adjusting the weight for WRR scheduling, the same concept is equally applicable to rate scheduling. By adjusting the cell-to-cell spacing, a class A connection can borrow bandwidth from a class B connection.

While the invention has been particularly shown and described by the foregoing detailed description, it will be understood by those skilled in the art that various changes, alterations, modifications, mutations and derivations in form and detail may be made without departing from the spirit and scope of the invention.

**WHAT IS CLAIMED IS:**

1. A method of adaptive resource allocation in transmitting data, comprising:  
allocating a resource to each of a plurality of data transmitting active connections, the plurality of  
active connections belonging to more than one class of service;  
5 determining a lender class of service for each active connection from which resources may be  
reallocated to the active connection;  
periodically comparing the resource usage of an active connection to an upper threshold and a  
lower threshold;  
borrowing resources from the lender class of service for the active connection in response to the  
10 current usage exceeding the upper threshold; and  
returning resources to the lender class of service from the active connection in response to the  
current usage being less than the lower threshold.

2. The method, as set forth in claim 1, wherein periodically comparing the resource usage  
15 comprises comparing the resource usage of the active connection in response to receiving incoming data  
on the active connection.

3. The method, as set forth in claim 1, wherein periodically comparing the resource usage  
comprises comparing the resource usage of the active connection in response to receiving an incoming  
20 cell of data on the active connection.

4. The method, as set forth in claim 1, further comprising:  
recording a time of the last resource reallocation action when resource was borrowed or returned  
for the active connection;  
25 comparing a current time with the last resource reallocation action time and proceeding or not  
proceeding to borrowing in response to the difference therebetween.

5. The method, as set forth in claim 1, further comprising:  
recording a time of the last resource reallocation action when resource was borrowed or returned  
30 for the active connection;  
comparing a current time with the last resource reallocation action time; and  
proceeding to borrowing in response to the difference between the current time and the last  
resource reallocation action time being greater than a predetermined minimum time spacing requirement,  
and not proceeding to borrowing in response to the difference between the current time and the last

resource reallocation action time being less than or equal to the predetermined minimum time spacing requirement.

6. The method, as set forth in claim 1, further comprising:

5 recording a time of an oldest borrow action when resource was borrowed for the active connection;

comparing a current time with the oldest borrow action time; and

proceeding to returning in response to the difference between the current time and the oldest borrow action time being greater than a predetermined maximum borrow time requirement.

10 7. The method, as set forth in claim 1, wherein borrowing comprises:

determining available resources of the lender class of service; and

proceed to borrowing if the available resources is sufficient.

15 8. The method, as set forth in claim 1, further comprising:

storing data associated with each borrow action for each active connection, including:

a borrow action time;

an amount of resource borrowed for the active connection; and

an weight increment for the active connection.

20 9. The method, as set forth in claim 1, further comprising:

storing data associated with each active connection, including:

a time of last resource reallocation for the active connection;

25 an enable indicative of whether resource reallocation is permitted for the active connection;

a lender class identifier indicative of the class of service that resources can be borrowed from;

an upper threshold indicative of need to borrow resources; and

a lower threshold indicative of need to return borrowed resources.

30 10. The method, as set forth in claim 9, wherein storing data associated with each active connection further comprises:

storing a minimum time spacing between resource reallocation actions; and

storing a maximum time to return borrowed resources.

11. The method, as set forth in claim 8, further comprising storing data associated with each borrow action for each active connection in a linked list, with the oldest borrow action at the head of the linked list.

5 12. The method, as set forth in claim 11, wherein returning borrowed resources comprises removing an oldest link in the linked list.

13. The method, as set forth in claim 1, further comprising first scheduling resources using a weighted round robin method.

10 14. The method, as set forth in claim 1, wherein borrowing and returning resources comprise borrowing and returning bandwidths.

15 15. The method, as set forth in claim 1, wherein borrowing and returning resources comprise increasing and decreasing data sending rate.

16. A telecommunications equipment having limited resources to allocate to active connections for a plurality of classes of services, using a method of dynamic resource allocation, comprising:

20 a first data structure storing resource allocation weighting data associated with each class of services;

a second data structure storing resource allocation weighting data associated with each active connection in each class of services;

wherein the second data structure includes:

25 a class identifier specifying a lender class of services from which resource may be borrowed;

an upper threshold indicative of when resource borrowing action for this active connection should be initiated;

30 a lower threshold indicative of when resource returning action for this active connection should be initiated; and

a sequence of data associated with each resource borrowing action of resources that have not been returned.

17. The telecommunications equipment, as set forth in claim 16, wherein the second data structure further comprises a timestamp of the most recent borrow or return action for the active connection.

5 18. The telecommunications equipment, as set forth in claim 16, wherein the second data structure further comprises a minimum time requirement between borrow and return actions for the active connection.

10 19. The telecommunications equipment, as set forth in claim 16, wherein the second data structure further comprises an amount of resource that may be borrowed by the active connection from the lender class of services in each borrow action.

15 20. The telecommunications equipment, as set forth in claim 16, wherein the second data structure further comprises a maximum time requirement that borrowed resources by the active connection need to be returned to the lender class of services.

21. The telecommunications equipment, as set forth in claim 16, wherein the sequence of data comprises a linked list of data associated with each borrow action.

20 22. The telecommunications equipment, as set forth in claim 21, wherein the data associated with each borrow action comprises a timestamp of the borrow action.

25 23. The telecommunications equipment, as set forth in claim 21, wherein the data associated with each borrow action comprises a weighting increment for the active connection, and a weight decrement of the lender class of services.

24. A method of adaptive dynamic weight assignment for transmitting data, comprising:  
assigning a total weight to each class of services;  
assigning a weight to each of a plurality of data transmitting active connections in each class of  
30 services;

determining a lender class of service for each active connection from which weighting may be reallocated to the active connection;

periodically comparing the resource usage of an active connection to an upper threshold and a lower threshold;



borrowing additional weighting from the lender class of service for the active connection in response to the current usage exceeding the upper threshold; and

returning borrowed weighting to the lender class of service from the active connection in response to the current usage being less than the lower threshold.

5

25. The method, as set forth in claim 24, wherein periodically comparing the resource usage comprises comparing the resource usage of the active connection in response to receiving incoming data on the active connection.

10

26. The method, as set forth in claim 24, wherein periodically comparing the resource usage comprises comparing the resource usage of the active connection in response to receiving an incoming cell of data on the active connection.

15

27. The method, as set forth in claim 24, further comprising:  
recording a time of the last weighting reallocation action when weighting was borrowed or returned for the active connection;

comparing a current time with the last weighting reallocation action time and proceeding or not proceeding to borrowing in response to the difference therebetween.

20

28. The method, as set forth in claim 24, further comprising:  
recording a time of the last weighting reallocation action when weighting was borrowed or returned for the active connection;

comparing a current time with the last weighting reallocation action time; and

proceeding to borrowing in response to the difference between the current time and the last weighting reallocation action time being greater than a predetermined minimum time spacing requirement, and not proceeding to borrowing in response to the difference between the current time and the last weighting reallocation action time being less than or equal to the predetermined minimum time spacing requirement.

30

29. The method, as set forth in claim 24, further comprising:  
recording a time of an oldest borrow action when weighting was borrowed for the active connection;

comparing a current time with the oldest borrow action time; and

proceeding to returning in response to the difference between the current time and the oldest borrow action time being greater than a predetermined maximum borrow time requirement.

35

30. The method, as set forth in claim 24, wherein borrowing comprises:  
determining available weighting of the lender class of service; and  
proceed to borrowing if the available weighting is sufficient.

5

31. The method, as set forth in claim 24, further comprising:  
storing data associated with each borrow action for each active connection, including:  
a borrow action time;  
an amount of resource borrowed for the active connection; and  
an weight increment for the active connection.

10

32. The method, as set forth in claim 24, further comprising:  
storing data associated with each active connection, including:  
a time of last weighting reallocation for the active connection;  
an enable indicative of whether weighting reallocation is permitted for the active  
connection;  
a lender class identifier indicative of the class of service that weighting can be borrowed  
from;  
an upper threshold indicative of need to borrow weighting; and  
a lower threshold indicative of need to return borrowed weighting.

15

20

33. The method, as set forth in claim 32, wherein storing data associated with each active  
connection further comprises:  
storing a minimum time spacing between resource reallocation actions; and  
storing a maximum time to return borrowed resources.

25




34. The method, as set forth in claim 31, further comprising storing data associated with each  
borrow action for each active connection in a linked list, with the oldest borrow action at the head of the  
linked list.

30

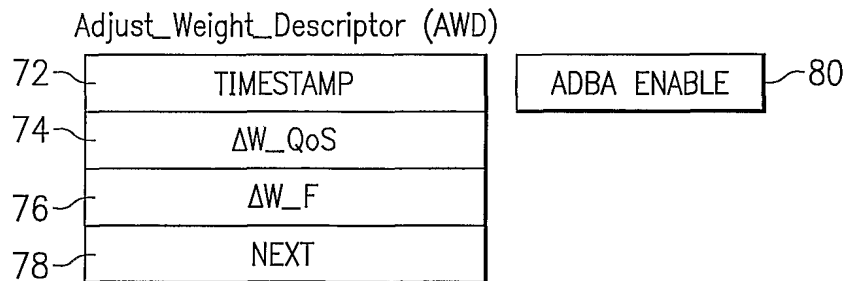

35. The method, as set forth in claim 34, wherein returning borrowed weighting comprises  
removing an oldest link in the linked list.

35

36. The method, as set forth in claim 24, further comprising using a weighted round robin  
method to schedule bandwidth assignment.

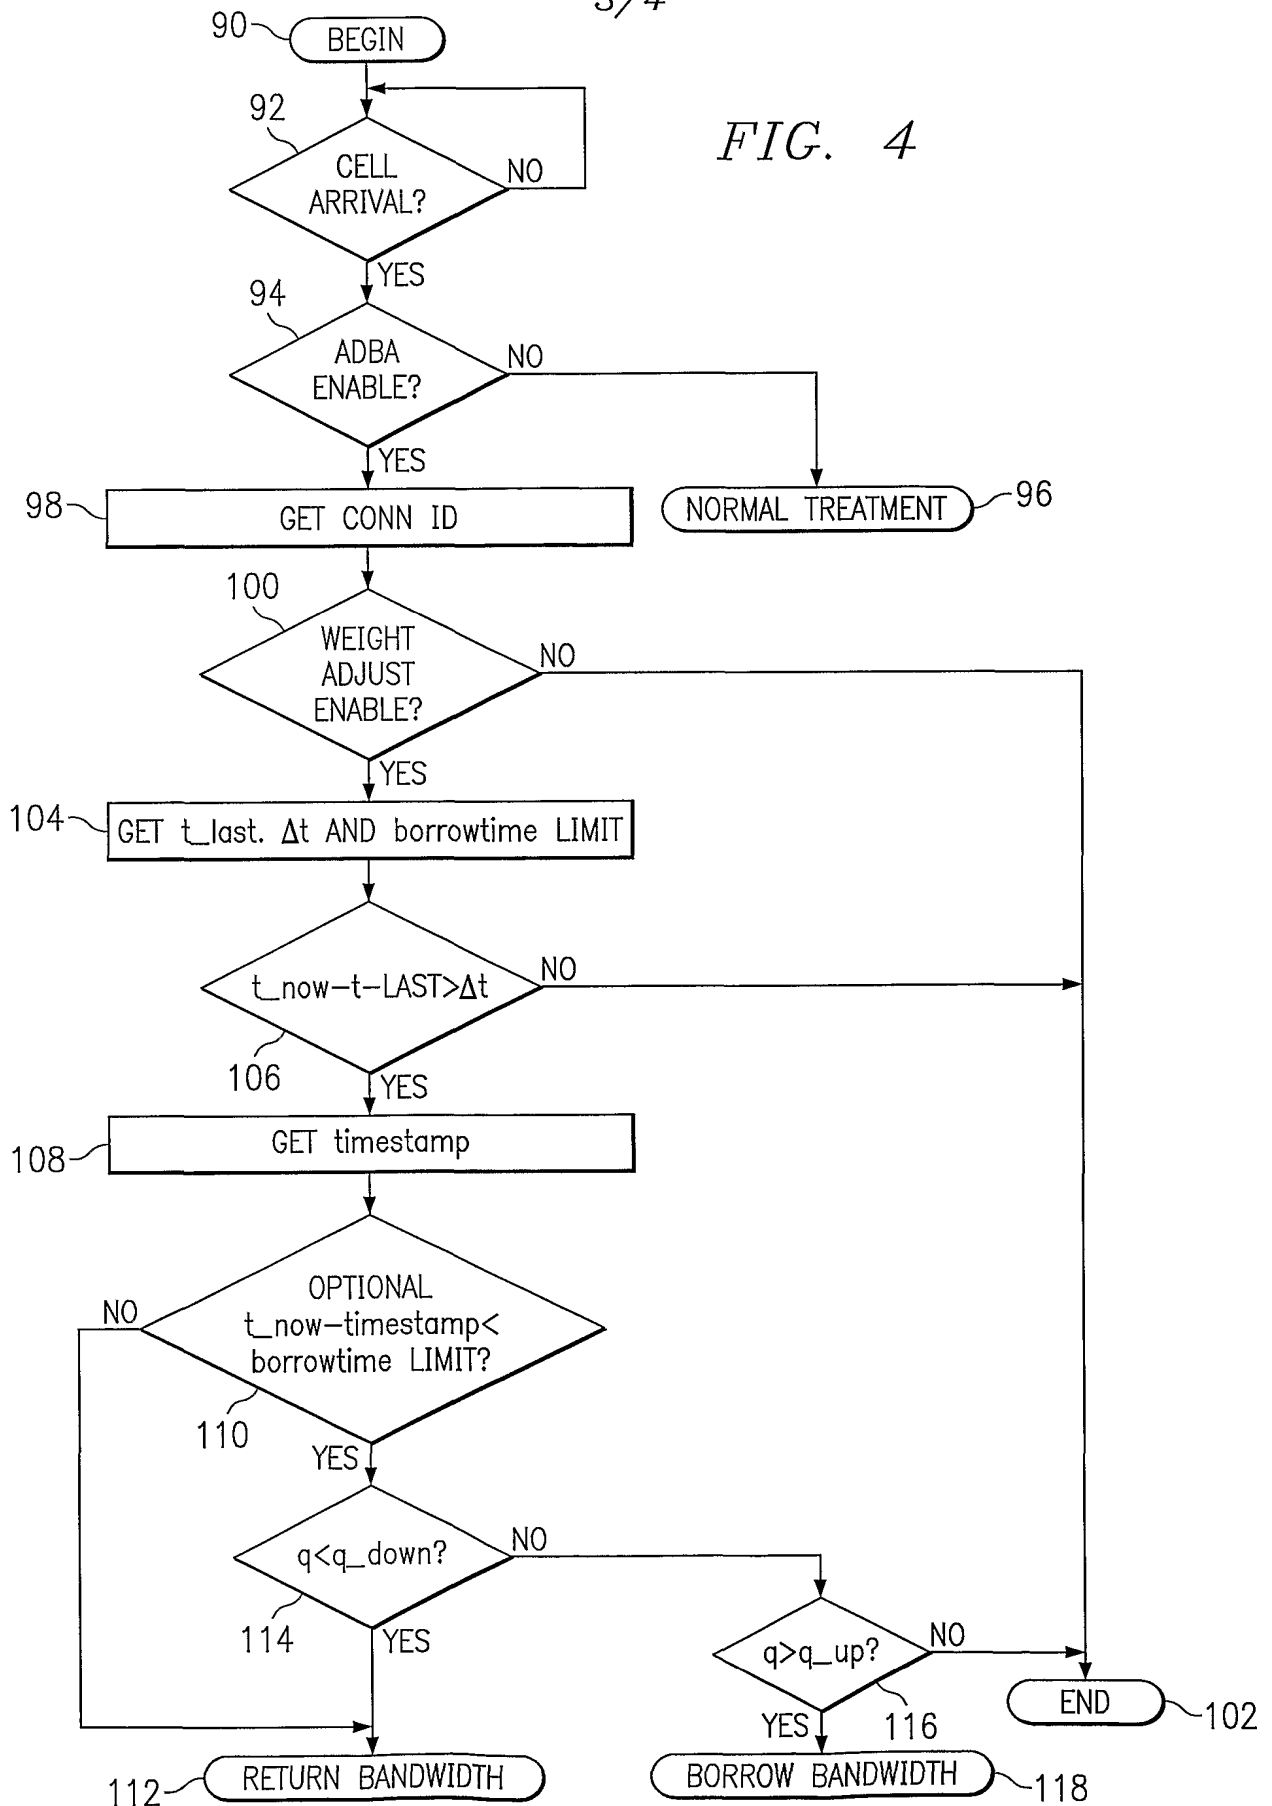


2/4

43  
*FIG. 3C**FIG. 5*

3/4

FIG. 4



4/4

FIG. 6

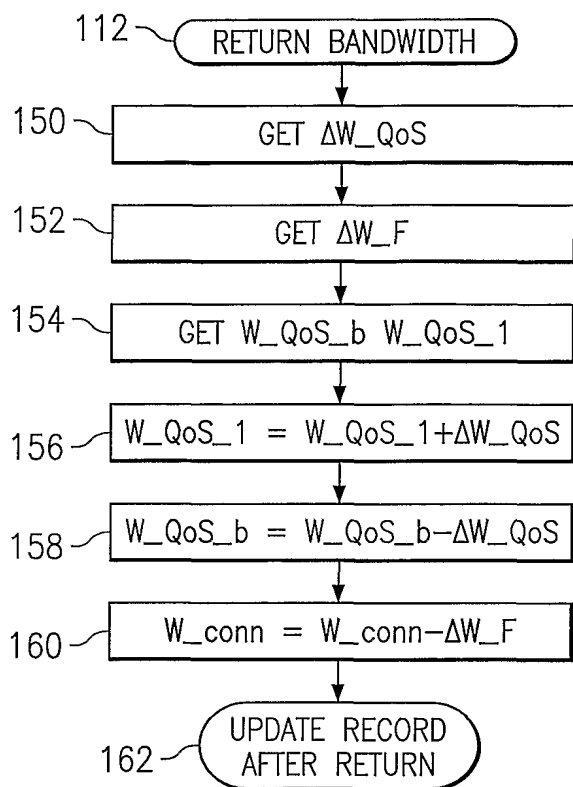


FIG. 7

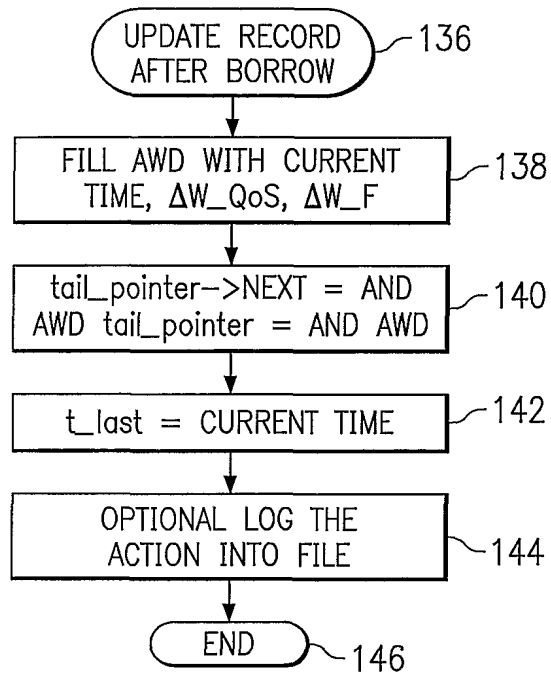


FIG. 8

